IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Title of the Invention

Cross-Selling Optimizer

Inventors

Marc-David Cohen

Judith Tyler Parks

Cross-Selling Optimizer

RELATED APPLICATION

This application claims priority to U.S. provisional application Serial No. 60/207,609 entitled CROSS SELLING OPTIMIZER filed May 26, 2000. By this reference, the full disclosure, including the drawings, of U.S. provisional application Serial No. 60/207,609 are incorporated herein.

BACKGROUND OF THE INVENTION

1. Technical Field

The present invention is generally directed to computer-implemented sales data analysis, and more specifically to computer-implemented marketing and selling efforts optimization.

2. Description of the Related Art

Previous Customer Relationship Management (CRM) solutions involve a combination of data warehousing and data mining. These components form the basis for collecting information for understanding customer relationships and potential market growth. Identifying cross-selling opportunities is an important goal of the CRM solution. One way this is done in CRM is with market basket analysis. Based on the principles of market basket analysis, the association node in a data miner (such as the data miner "Enterprise Miner" available from SAS Institute Inc.) produces rules data that show cross-selling opportunities. However, the rules do not show which of these opportunities is best in meeting overall business goals nor do they show how to distribute resources to achieve those business goals.

For example, the rules data may contain an association rule such as “CKING → SVG & CCRD” with a statistical "lift" value of 1.1. This is often interpreted to mean that the population that only has purchased a check product ("CKING") has some potential, of strength 1.1, to purchase savings accounts and credit card products (respectively, "SVG" and "CCRD"). While of value in identifying specific customer populations’ potential, the solution gives no suggestion as to whether this or any other rule should be used as a basis for expansion of just the savings account market. Moreover, if this rule is used as a basis for allocating resources it does not show how that decision will impact the target for the CCRD market and whether exploiting this rule is a good overall use of limited resources. Thus, the present approach has difficulty in addressing such business problems as: How can I best allocate limited resources to exploit cross-selling opportunities that meet my overall product sales goals?

SUMMARY OF THE INVENTION

A cross-selling optimization (CSO) method and system are provided for allocating marketing and selling effort in the cross-selling environment. It addresses the problem of optimizing cross-selling efforts as well as other problems in the previous approaches. It optimally allocates resources based on results from data warehousing and data mining methodologies. These methodologies form the basis for collecting information for understanding customer relationships and potential market growth. The present invention preferably uses linear programming to determine the optimal way in which to allocate limited cross-selling resources to marketing various products so that the highest possible return on one’s marketing investment (ROI) is achieved.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention satisfies the general needs noted above and provides many advantages, as will become apparent from the following description when read in conjunction with the accompanying drawings, wherein:

FIG. 1 is a block diagram depicting the module structure and data flow of the present invention;

FIG. 2 is a table depicting an exemplary association rule dataset;

FIG. 3 is a table depicting an exemplary subset association rules dataset as generated by the present invention;

FIG. 4 is a pie chart that graphically depicts the exemplary subset association rules dataset of FIG. 3;

FIG. 5 is a table depicting how the effort applied to the targeted populations of FIG. 3 translates into effort applied to products; and

FIG. 6 is a bar chart that graphically depicts the tabular values of FIG. 5.

DETAILED DESCRIPTION

FIG. 1 depicts the cross-selling optimization system of the present invention as generally shown by reference numeral 20. The cross-selling optimization system 20 generates subset association rules 38 based upon raw data 22 that has been pre-processed by a data miner 24. The cross-selling optimization system 20 includes an optimization model 32 (e.g., a linear programming model) to generate the subset association rules 38.

The subset association rules dataset 38 addresses one or more business issues by associating cross-selling rules with metrics that solve the business issues. For example, a

business issue may address where should a business allocate its limited personnel efforts so as to optimize its overall return on investment. The present invention provides in subset association rules dataset 38 what amount of effort should be allocated to each cross-selling opportunity so as to optimize the overall return on investment.

First, the present invention may use a data miner 24 to process raw data 22. Raw data 22 may include historical data on the product sales of the business. An exemplary data miner 22 is Enterprise Miner™ available from the SAS Institute Inc. of Cary, North Carolina. Enterprise Miner™ processes the raw data according to known techniques to generate association rules data set 26.

The association rules dataset 26 contains variables such as a rule and at least one cross-selling statistic as shown by box 28. The rules variable, for example, lists products on the left-hand-side of the arrow and products on the right-hand-side. Thus, each record in the dataset is a new rule that elucidates a unique customer group or segment. In addition, each record provides at least one cross-selling statistic to convey information on the likelihood of selling the right-hand-side products to customers who have the left-hand-side products. These statistics or metrics are calculated from raw data 22.

Statistics for this may include the lift and the expected confidence. The lift is the ratio of the probability of having the right-hand-side product(s) given that the customer has the left-hand-side product(s), over the probability that the customer has the right-hand-side product(s). Thus, a large value of lift indicates that the percentage of population with the left-hand-side product(s) is relatively small compared to the strength of the relationship between the right-hand-side and left-hand-side product(s). Other cross-selling statistical metrics may

be used in combination with the lift variable to convey additional information on a cross-selling likelihood. For example, the E_Confidence variable may be used with the lift variable to indicate the frequency with which the right-hand-side product occurs in the overall population.

FIG. 2 shows a sample of the association rules dataset 26 where each row reflects a different product combination. For example, row 50 shows that if a person buys a saving accounts product, then the person is likely to purchase credit card products since the lift value is greater than 1. For the products listed in FIG. 2, the following list provides what product a symbol denotes:

Symbol	**Product**
ATM	Automatic teller machine
AUTO	Auto loan
CCRD	Credit card
CD	Certification of deposit
CKCRD	Check card
CKING	Checking Account
HMEQLC	Home equity loan
IRA	Individual retirement account
MMDA	Money market certificate
SVG	Savings Account

The raw dataset 22 is analyzed by the association's node in Enterprise Miner™, which generates association rules dataset 26. One technique to determine these rules is by counting the number of customers in the database that have the different combinations of products. Analysts may use these totals to make inferences about the likelihood of successfully selling new products to existing customers. In this way the rules identify cross-selling opportunities. However, the rules do not show which of these opportunities is best in meeting one or more overall business goals, nor do they show how to distribute resources to achieve those business goals while maintaining a high return on investment.

The present invention addresses these problems by capturing business issues 34 in an optimization model 32. Business issues 34 may be external resource goals or effort targets for each individual product. Resources can be measured in many different ways such as dollars, people, and person-hours. To model or represent this effort resource, one assumes that there is one (1) unit resource available (this is 100% selling effort). One also assumes that there are target effort levels for each individual product. That is, a target percentage of effort to spend on selling each product is known.

Box 30 represents the construction of an optimization model 32 based upon the captured business issues of box 34. The optimization model 32 may be a linear program model or some other type of optimization program, such as a non-linear optimization program. The model includes a business objective function and a set of business constraints as shown by reference numeral 31. The business constraints may be stored in data structures that are accessible through any conventional computer storage memory devices.

The business objective function drives the calculation of optimal amount of effort, and the constraints capture the various business issues 34. For example, an objective may be to distribute 100% selling effort across the multitude of cross-selling opportunities so that the weighted average of the product of lift and potential revenue is maximized. This average is weighted by the effort. The optimization model 32 may also have as input: user supplied parameters such as the product effort target levels, the anticipated returns from selling to different customer groups, and maximum acceptable average expected confidence.

Box 36 shows the model being solved using SAS/OR®. SAS/OR® is a comprehensive set of enterprise decision-making tools available from the SAS Institute Inc. The solution is a vector of efforts to be applied to the different customer groups. More specifically, the solution is the subset association rules dataset 38. Dataset 38 contains a subset of association rules together with optimal amounts of effort resources to expend.

An example of the subset association rules dataset 38 is shown in FIG. 3. Each row of the dataset indicates a product on which it would be useful for the business to expend its cross-selling resources. The different columns show various unique customer groups or segments (represented by unique combinations of products), optimal amounts of resource to expend marketing each individual product, and the actual product(s) to be marketed. For example, row 60 includes the cross-selling opportunities metric of "Effort" to illustrate that it is optimal to allocate 10% of one's resources (e.g., 10% of the marketing budget) towards marketing savings accounts and CD products to the customer segment that holds only checking accounts.

In building the model at box 30, certain assumptions are defined. As an example, we may assume that there is a limited amount of resource to expend on product marketing and selling. This is termed the resource effort, and it is assumed that there is one (1) unit available (this is 100% selling effort). It is also assumed that there are target effort levels for each individual product. That is, we know a target percentage of effort to spend on selling each product. Note that effort can be measured in many different kinds of units. It could be dollars, people, or person hours. For example, you as a marketing director may have a budget of $4 million dollars to spend on product marketing and selling. This $4 million dollars represents 100% effort resource. In addition, you may have a specific product effort target in that you want to spend $500,000 marketing IRAs.

Thus, the business problem in this example is posed as, on which customer groups should you focus your selling efforts in order to meet your targets for each product and, at the same time, maximize the return on your effort investment? This is called the objective. The solution answers this by identifying the amount of effort to use on each customer group. The solution also meets the product sales targets while maximizing the return on the investment.

Information about the size of the potential markets is incorporated implicitly in the objective through the lift. Since this is accounted for implicitly, the solution may recommend significant effort for customer groups simply because they have a large likelihood of success even though they do not represent a large market. To provide some control on this, a constraint is added that limits the average expected confidence weighted by effort to be less than a user supplied quantity.

In this example, there are three types of constraints. One constraint specifies that the total amount of effort is 1. This defines the limited resources available for selling. Another restricts the average of expected confidence weighted by effort. This biases the effort towards customer populations that have greater growth potential. Finally, there is a set of constraints that requires a certain amount of effort be allocated to each product.

The model can be specified unambiguously as follows. Let

$J =$ set of products j

$I =$ set of rules i

$$a_{ij} = \begin{cases} 1 & \text{if rule } i \text{ has product } j \text{ as a result} \\ 0 & \end{cases}$$

$T_j =$ target effort for product j

$r_i =$ return from rule i if 100% effort is applied

$l_i =$ lift from rule i

$c_i =$ expected confidence of rule i

$C =$ maximum expected confidence for the weighted average effort allocation

$x_i =$ effort to apply to rule i

All of these quantities are known input parameters except for the effort x_i. This is the quantity that is to be calculated by the system and is called the decision variable.

Formal specification has the objective as

$$\text{Max} \sum_{i \in I} r_i l_i x_i$$

and the constraints as:

$$\sum_{i \in I} a_{ij} x_i \geq T_j \quad \forall j \in J \qquad \text{Target product efforts}$$

$$\sum_{i \in I} x_i = 1 \qquad \text{100\% effort available}$$

$$\sum_{i \in I} c_i x_i \leq C \qquad \text{Confidence limit}$$

$$x_i \geq 0 \quad \forall i \in I \qquad \text{Nonnegative effort}$$

This example has assumed that the product targets, T_j, are identical for all the products and that the returns, r_i, are 1 for each rule in the data set. The present invention uses a software macro that has three arguments: "ds=" which is the name of the rules data set; "conf=" which is the value for the limit on the weighted average expected confidence; and "target=" which is the target effort level for each of the products. The macro call looks like:

```
%CSO(ds=,conf=,target=);
```

The following example illustrates the present invention. The macro is called with an expected confidence level of 25 and an effort target of 10% for each of the 10 products that appear on the right-hand-side of at least one rule in the data set.

```
%CSO(ds=crm.rules,conf=25,target=.1);
```

The macro solves the problem by finding the distribution of effort that meets the constraints discussed above and maximizes the total lift weighted by effort (since the returns r_i are all 1) using known linear programming techniques.

FIG. 3 depicts in a table format the solution. The present invention has selected those customer populations that should be marketed or sold to. It shows the amount of effort to be applied to each of these targeted populations. The present invention picked populations that

tend to have larger lift as we would expect, because of the objective. Also, it should be noted that most of the populations have effort .1, except for **"CKING & CCRD → CKCRD"** which has effort .4. Most likely this group is selected because of its high lift and low expected confidence. In general, the present invention picks populations that have small expected confidence because of the constraint limiting the weighted average expected confidence to 25. FIG. 4 depicts graphically in a pie chart format the tabular results of FIG. 3.

FIG. 5 depicts in a tabular format how the effort applied to these targeted populations translates into effort applied to products. Note that each product has at least .1 effort as is required by the business objectives. Note that the CKING product has a total effort of 0.5 due to its combined values of 0.4 and 0.1 as shown respectively at rows 70 and 72. FIG. 6 is a bar chart representation of FIG. 5.

The preferred embodiment described with reference to the drawing figures and associated tables is presented only to demonstrate examples of the present invention. Additional and/or alternative embodiments of the present invention should be apparent to one of ordinary skill in the art upon reading this disclosure. For example, alternative elements or steps that may be included in the present invention include: a constraint that seeks to ensure even dispersion of effort throughout all products; and a constraint that ensures a respectable return on equity. These constraints address such additional business issues as an organization being more interested in maintaining a certain level of shareholder value by avoiding inadequately performing products, rather than maintaining a diverse market of products. As a further example of the broad range of alternate embodiments, a performance measure other

than the lift may be used as a measure of potential of the customer population. This may include using the lift factor divided by the maximum lift over all products, as a relative measure of potential.

The present invention also can be used to analyze cross-selling efforts on a regional basis. In this situation, the association rules and statistics would include geographical information in order to determine what are the optimal effort allocations on a per region basis. Still further, the present invention analyzes cross-selling efforts involving items other than products, such as the sale of services.